

Efficient Sparse Group Feature Selection via Nonconvex Optimization

Shuo Xiang^{1,2} Xiaotong Shen³ Jieping Ye^{1,2}

¹Department of Computer Science and Engineering, Arizona State University, AZ 85287

²Center for Evolutionary Medicine and Informatics, The Biodesign Institute, Arizona State University, AZ 85287

³School of Statistics, University of Minnesota, Minneapolis, MN 55455

January 22, 2013

Abstract

Sparse feature selection has been demonstrated to be effective in handling high-dimensional data. While promising, most of the existing works use convex methods, which may be suboptimal in terms of the accuracy of feature selection and parameter estimation. In this paper, we expand a nonconvex paradigm to sparse group feature selection, which is motivated by applications that require identifying the underlying group structure and performing feature selection simultaneously. The main contributions of this article are twofold: (1) statistically, we introduce a nonconvex sparse group feature selection model which can reconstruct the oracle estimator. Therefore, consistent feature selection and parameter estimation can be achieved; (2) computationally, we propose an efficient algorithm that is applicable to large-scale problems. Numerical results suggest that the proposed nonconvex method compares favorably against its competitors on synthetic data and real-world applications, thus achieving desired goal of delivering high performance.

1 Introduction

During the past decade, sparse feature selection has been extensively investigated, on both optimization algorithms [1] and statistical properties [28, 20, 3]. When the data possesses certain group structure, sparse modeling has been explored in [24, 16, 13] for group feature selection. The group lasso [24] proposes an L_2 -regularization method for each group, which ultimately yields a group-wisely sparse model. The utility of such a method has been demonstrated in detecting splice sites [23]—an important step in gene finding and theoretically justified in [13]. The sparse group lasso [11] enables to encourage sparsity at the level of both features and groups simultaneously.

In the literature, most approaches use convex methods due to globality of the solution and tractable computation. However, this may lead to suboptimal results. Recent studies demonstrate that nonconvex methods, for instance, the truncated L_1 -penalty [19, 15, 27], may have potential to deliver superior performance than the standard L_1 -formulation. In addition, [19] suggests that a constrained version of nonconvex regularization is slightly more preferable than its regularization counterpart due to theoretical merits.

In this article, we investigate the sparse group feature selection (SGFS) through a constrained nonconvex

formulation. Ideally, we wish to optimize the following L_0 -model:

$$\begin{aligned}
& \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \\
& \text{subject to} && \sum_{j=1}^p I(|x_j| \neq 0) \leq s_1 \\
& && \sum_{j=1}^{|G|} I(\|\mathbf{x}_{G_j}\|_2 \neq 0) \leq s_2,
\end{aligned} \tag{1}$$

where \mathbf{A} is an n by p data matrix with its columns representing different features. $\mathbf{x} = (x_1, \dots, x_p)$ is partitioned into $|G|$ non-overlapping groups $\{\mathbf{x}_{G_i}\}$ and $I(\cdot)$ is the indicator function. The advantage of the L_0 -model (1) lies in its complete control on two levels of sparsity (s_1, s_2), which are the numbers of features and groups respectively. However, a problem like (1) is known to be NP-hard [17].

This paper develops an efficient nonconvex method, which is a computational surrogate of the L_0 -method described above and has theoretically guaranteed performance. We contribute in two aspects: (i) statistically, the proposed method retains the merits of the L_0 approach (1) in the sense that the oracle estimator can be reconstructed, which leads to consistent feature selection and parameter estimation; (ii) computationally, our efficient optimization tool enables to treat large-scale problems.

2 Nonconvex Formulation and Computation

One major difficulty of solving (1) comes from nonconvex and discrete constraints, which require enumerating all possible combinations of features and groups to achieve the optimal solution. Therefore we approximate these constraints by their continuous computational surrogates:

$$\begin{aligned}
& \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \\
& \text{subject to} && \sum_{j=1}^p J_\tau(|x_j|) \leq s_1, \\
& && \sum_{i=1}^{|G|} J_\tau(\|\mathbf{x}_{G_i}\|_2) \leq s_2,
\end{aligned} \tag{2}$$

where $J_\tau(z) = \min(|z|/\tau, 1)$ is a truncated L_1 -function approximating the L_0 -function [19, 26], and $\tau > 0$ is a tuning parameter such that $J_\tau(z)$ approximates the indicator function $I(|z| \neq 0)$ as τ approaches zero.

To solve the nonconvex problem (2), we develop a Difference of Convex (DC) algorithm based on a decomposition of each nonconvex constraint function into a difference of two convex functions; for instance,

$$\sum_{j=1}^p J_\tau(|x_j|) = S_1(\mathbf{x}) - S_2(\mathbf{x}),$$

where

$$S_1(\mathbf{x}) = \frac{1}{\tau} \sum_{j=1}^p |x_j|$$

and

$$S_2(\mathbf{x}) = \frac{1}{\tau} \sum_{j=1}^p \max\{|x_j| - \tau, 0\}$$

are convex in \mathbf{x} . Then each trailing convex function, say $S_2(\mathbf{x})$, is replaced by its affine minorant at the previous iteration

$$S_1(\mathbf{x}) - S_2(\hat{\mathbf{x}}^{(m-1)}) - \nabla S_2(\hat{\mathbf{x}}^{(m-1)})^T(\mathbf{x} - \hat{\mathbf{x}}^{(m-1)}), \quad (3)$$

which yields an upper approximation of the constraint function $\sum_{j=1}^p J_\tau(|x_j|)$ as follows:

$$\frac{1}{\tau} \sum_{j=1}^p |x_j| \cdot I(|\hat{x}_j^{(m-1)}| \leq \tau) + \sum_{j=1}^p I(|\hat{x}_j^{(m-1)}| > \tau) \leq s_1. \quad (4)$$

Similarly, the second nonconvex constraint in (2) can be approximated by

$$\frac{1}{\tau} \sum_{j=1}^{|G|} \|\mathbf{x}_{G_j}\|_2 \cdot I(\|\hat{\mathbf{x}}_{G_j}^{(m-1)}\|_2 \leq \tau) + \sum_{j=1}^{|G|} I(\|\hat{\mathbf{x}}_{G_j}^{(m-1)}\|_2 > \tau) \leq s_2. \quad (5)$$

Note that both (4) and (5) are convex constraints, which result in a convex subproblem as follows:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \\ & \text{subject to} && \frac{1}{\tau} \|\mathbf{x}^{T_1(\hat{\mathbf{x}}^{(m-1)})}\|_1 \leq s_1 - (p - |T_1(\hat{\mathbf{x}}^{(m-1)})|) \\ & && \frac{1}{\tau} \|\mathbf{x}^{T_3(\hat{\mathbf{x}}^{(m-1)})}\|_G \leq s_2 - (|G| - |T_2(\hat{\mathbf{x}}^{(m-1)})|), \end{aligned} \quad (6)$$

where T_1 , T_2 and T_3 are the support sets¹ defined as:

$$\begin{aligned} T_1(\mathbf{x}) &= \{i : |x_i| \leq \tau\} \\ T_2(\mathbf{x}) &= \{i : \|\mathbf{x}_{G_i}\|_2 \leq \tau\} \\ T_3(\mathbf{x}) &= \{i : x_i \in \mathbf{x}_{G_j}, j \in T_2(\mathbf{x})\}, \end{aligned}$$

$\|\mathbf{x}^{T_1}\|_1$ and $\|\mathbf{x}^{T_3}\|_G$ denote the corresponding value restricted on T_1 and T_3 respectively, and $\|\mathbf{x}\|_G = \sum_{i=1}^{|G|} \|\mathbf{x}_{G_i}\|_2$. Solving (6) would provide us an updated solution, denoted as $\hat{\mathbf{x}}^{(m)}$. Such procedure is iterated until the objective value is no longer decreasing, indicating that a local minimizer is achieved. The DC algorithm is summarized in Algorithm 1, from which we can see that efficient computation of (6) is critical to the overall DC routine. We defer detailed discussion of this part to Section 4.

Algorithm 1 DC programming for solving (2)

Input: \mathbf{A} , \mathbf{y} , s_1 , s_2

Output: solution \mathbf{x} to (2)

- 1: **(Initialization)** Initialize $\hat{\mathbf{x}}^{(0)}$.
 - 2: **(Iteration)** At iteration m , compute $\hat{\mathbf{x}}^{(m)}$ by optimizing (6).
 - 3: **(Stopping Criterion)** Terminate when the objective function stops decreasing.
-

3 Theoretical Results

This section investigates theoretical aspects of the proposed method. More specifically, we demonstrate that the oracle estimator $\hat{\mathbf{x}}^o$, the least squares estimator based on the true model, can be reconstructed. As a result, consistent selection as well as optimal parameter estimation can be achieved.

¹Support sets indicate that the elements outside these sets have no effect on the particular items in the constraints of (6).

For better presentation, we introduce some notations that would be only utilized in this section. Let $C = (G_{i_1}, \dots, G_{i_k})$ be the collection of groups that contain nonzero elements. Let $A_{G_j} = A_{G_j}(\mathbf{x})$ and $A = A(\mathbf{x})$ denote the indices of nonzero elements of \mathbf{x} in group G_j and in entire \mathbf{x} respectively. Define

$$\mathcal{S}_{j,i} = \{\mathbf{x} \in \mathcal{S} : (A_C, C) \neq (A_{C^0}, C^0), |A| = j, |C| = i\},$$

where \mathcal{S} is the feasible region of (2) and C^0 represents the true nonzero groups.

The following assumptions are needed for obtaining consistent reconstruction of the oracle estimator:

Assumption 1 (Separation condition). *Define*

$$C_{\min}(\mathbf{x}^0) = \inf_{\mathbf{x} \in \mathcal{S}} \frac{-\log(1 - h^2(\mathbf{x}, \mathbf{x}^0))}{\max(|C^0 \setminus C|, 1)},$$

then for some constant $c_1 > 0$,

$$C_{\min}(\mathbf{x}^0) \geq c_1 \frac{\log |G| + \log s_1^0}{n},$$

where

$$h(\mathbf{x}, \mathbf{x}^0) = \left(\frac{1}{2} \int (g^{1/2}(\mathbf{x}, y) - g^{1/2}(\mathbf{x}^0, y))^2 d\mu(y) \right)^{1/2}$$

is the Hellinger-distance for densities with respect to a dominating measure μ .

Assumption 2 (Complexity of the parameter space). *For some constants $c_0 > 0$ and any $0 < t < \varepsilon \leq 1$,*

$$H(t, \mathcal{F}_{j,i}) \leq c_0 \max((\log(|G| + s_1^0))^2, 1) |\mathcal{B}_{j,i}| \log(2\varepsilon/t),$$

where $\mathcal{B}_{j,i} = \mathcal{S}_{j,i} \cap \{\mathbf{x} \in h(\mathbf{x}, \mathbf{x}^0) \leq 2\varepsilon\}$ is a local parameter space and $\mathcal{F}_{j,i} = \{g^{1/2}(\mathbf{x}, y) : \mathbf{x} \in \mathcal{B}_{j,i}\}$ is a collection of square-root densities. $H(\cdot, \mathcal{F})$ is the bracketing Hellinger metric entropy of space \mathcal{F} [14].

Assumption 3. *For some positive constants d_1, d_2, d_3 with $d_1 > 10$,*

$$-\log(1 - h^2(\mathbf{x}, \mathbf{x}^0)) \geq -d_1 \log(1 - h^2(\mathbf{x}^\tau, \mathbf{x}^0)) - d_3 \tau^{d_2} p,$$

where $\mathbf{x}^\tau = (x_1 I(|x_1| \geq \tau), \dots, x_p I(|x_p| \geq \tau))$.

With the above assumptions hold, we can conclude the following non-asymptotic probability error bound regarding the reconstruction of the oracle estimator $\hat{\mathbf{x}}^o$.

Theorem 1. *Suppose that Assumptions 2 and 3 hold. For a global minimizer of (2) $\hat{\mathbf{x}}$ with $(s_1, s_2) = (s_1^0, s_2^0)$ and $\tau \leq \left(\frac{(d_1 - 10)C_{\min}(\mathbf{x}^0)}{d_3 d} \right)^{1/d_2}$, the following result hold:*

$$\mathbb{P}(\hat{\mathbf{x}} \neq \hat{\mathbf{x}}^o) \leq \exp \left(-c_2 n C_{\min}(\mathbf{x}^0) + 2(\log |G| + \log s_1^0) \right).$$

Moreover, with Assumption 1 hold, $\mathbb{P}(\hat{\mathbf{x}} = \hat{\mathbf{x}}^o) \rightarrow 1$ and

$$Eh^2(\hat{\mathbf{x}}, \mathbf{x}^o) = (1 + o(1)) \max(Eh^2(\hat{\mathbf{x}}^o, \mathbf{x}^0), \frac{s_1^0}{n})$$

as $n \rightarrow \infty$, $|G| \rightarrow \infty$.

Theorem 1 states that the oracle estimator $\hat{\mathbf{x}}^o$ can be accurately reconstructed, which in turn yields feature selection consistency as well as the recovery of the performance of the oracle estimator in parameter estimation. Moreover, according to Assumption 1, such conclusion still holds when $s_1^0 |G|$ grows in the order of $\exp(c_1^{-1} n C_{\min})$. This is in contrast to existing conclusions on consistent feature selection, where the number of candidate features should be no larger than $\exp(c^* n)$ for some c^* [28]. In this sense, the number

of candidate features is allowed to be much larger when an additional group structure is incorporated, particularly when each group contains considerable redundant features.

To our knowledge, our theory for the grouped selection is the first of this kind. However, it has a root in feature selection. The large deviation approach used here is applicable to derive bounds for feature selection consistency. In such a situation, the result agrees with the necessary condition for feature selection consistency for any method, except for the constants independent of the sample size [19]. In other words, the required conditions are weaker than those for L_1 -regularization [21]. The use of the Hellinger-distance is mainly to avoid specifying a sub-Gaussian tail of the random error. This means that the result continues to hold even when the error does not have a sub-Gaussian tail.

4 Optimization Procedures

As mentioned in Section 2, efficient computation of the convex subproblem (6) is of critical importance for the proposed DC algorithm. Note that (6) has an identical form of the constrained sparse group lasso problem:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \\ & \text{subject to} && \|\mathbf{x}\|_1 \leq s_1 \\ & && \|\mathbf{x}\|_G \leq s_2 \end{aligned} \tag{7}$$

except that \mathbf{x} is restricted to the two support sets. As to be shown in Section 4.3, an algorithm for solving (6) can be obtained through only a few modifications on that of (7). Therefore, we first focus on solving (7).

4.1 Accelerated Gradient Method

For large-scale problems, the dimensionality of data can be very high, therefore first-order optimization is often preferred. We adapt the well-known accelerated gradient method (AGM) [18, 2], which is commonly used due to its fast convergence rate.

To apply AGM to our formulation (7), the crucial step is to solve the following Sparse Group Lasso Projection (SGLP):

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \\ & \text{subject to} && \|\mathbf{x}\|_1 \leq s_1 && (C_1) \\ & && \|\mathbf{x}\|_G \leq s_2 && (C_2), \end{aligned} \tag{8}$$

which is an Euclidean projection onto a convex set and a special case of (7) when A is the identity. For convenience, let C_1 and C_2 denote the above two constraints in what follows.

Since the AGM is a standard framework whose efficiency mainly depends on that of the projection step, we leave the detailed description of AGM in the supplement and introduce the efficient algorithm for this projection step (8).

4.2 Efficient Projection

We begin with some special cases of (8). If only C_1 exists, (8) becomes the well-known L_1 -ball projection [9], whose optimal solution is denoted as $\mathcal{P}_1^{s_1}(\mathbf{v})$, standing for the projection of \mathbf{v} onto the L_1 -ball with radius s_1 . On the other hand, if only C_2 is involved, it becomes the group lasso projection, denoted as $\mathcal{P}_G^{s_2}$. Moreover, we say a constraint is *active*, if and only if an equality holds at the optimal solution x^* ; otherwise, it is *inactive*.

Preliminary results are summarized in Lemma 1:

Lemma 1. Denote a global minimizer of (8) as \mathbf{x}^* . Then the following results hold:

1. If both C_1 and C_2 are inactive, then $\mathbf{x}^* = \mathbf{v}$.
2. If C_1 is the only active constraint, i.e., $\|\mathbf{x}^*\|_1 = s_1$, $\|\mathbf{x}^*\|_G < s_2$, then $\mathbf{x}^* = \mathcal{P}_1^{s_1}(\mathbf{v})$
3. If C_2 is the only active constraint, i.e., $\|\mathbf{x}^*\|_1 < s_1$, $\|\mathbf{x}^*\|_G = s_2$, then $\mathbf{x}^* = \mathcal{P}_G^{s_2}(\mathbf{v})$

4.2.1 Computing \mathbf{x}^* from the optimal dual variables

Lemma 1 describes a global minimizer when either constraint is inactive. Next we consider the case in which both C_1 and C_2 are active. By the convex duality theory [6], there exist unique non-negative dual variables λ^* and η^* such that \mathbf{x}^* is also the global minimizer of the following regularized problem:

$$\underset{\mathbf{x}}{\text{minimize}} \quad \frac{1}{2}\|\mathbf{x} - \mathbf{v}\|_2^2 + \lambda^*\|\mathbf{x}\|_1 + \eta^*\|\mathbf{x}\|_G, \quad (9)$$

whose solution is given by the following Theorem.

Theorem 2 ([11]). The optimal solution \mathbf{x}^* of (9) is given by

$$\mathbf{x}_{G_i}^* = \max\{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2 - \eta^*, 0\} \frac{\mathbf{v}_{G_i}^{\lambda^*}}{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2} \quad i = 1, 2, \dots, |G| \quad (10)$$

where $\mathbf{v}_{G_i}^{\lambda^*}$ is computed via soft-thresholding [8] \mathbf{v}_{G_i} with threshold λ^* as follows:

$$\mathbf{v}_{G_i}^{\lambda^*} = \text{SGN}(\mathbf{v}_{G_i}) \cdot \max\{|\mathbf{v}_{G_i}| - \lambda^*, 0\},$$

where $\text{SGN}(\cdot)$ is the sign function and all the operations are taken element-wisely.

Theorem 2 gives an analytical solution of \mathbf{x}^* in an ideal situation when the values of λ^* and η^* are given. Unfortunately, this is not the case and the values of λ^* and η^* need to be computed directly from (8). Based on Theorem 2, we have the following conclusion characterizing the relations between the dual variables:

Corollary 1. The following equations hold:

$$\|\mathbf{x}^*\|_1 = \sum_{i=1}^{|G|} \max\{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2 - \eta^*, 0\} \frac{\|\mathbf{v}_{G_i}^{\lambda^*}\|_1}{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2} = s_1 \quad (11)$$

$$\|\mathbf{x}^*\|_G = \sum_{i=1}^{|G|} \max\{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2 - \eta^*, 0\} = s_2 \quad . \quad (12)$$

Suppose λ^* is given, then computing η^* from (12) amounts to solving a median finding problem, which can be done in linear time [9].

Finally, we treat the case of unknown λ^* (thus unknown η^*). We propose an efficient bisection approach to compute it.

4.2.2 Computing λ^* : bisection

Given an initial guess (estimator) of λ^* , says $\hat{\lambda}$, one may perform bisection to locate the optimal λ^* , provided that there exists an oracle procedure indicating if the optimal value is greater than $\hat{\lambda}^2$. This bisection method can estimate λ^* in logarithm time. Next, we shall design an oracle procedure.

²An upper bound and a lower bound of λ^* should be provided in order to perform the bisection. These bounds can be easily derived from the assumption that both C_1 and C_2 are active.

Let the triples

$$(\mathbf{x}^*, \lambda^*, \eta^*) = \text{SGLP}(\mathbf{v}, s_1, s_2)$$

be the optimal solution of (8) with both constraints active, i.e., $\|\mathbf{x}^*\|_1 = s_1$, $\|\mathbf{x}^*\|_G = s_2$, with (λ^*, η^*) be the optimal dual variables. Consider the following two sparse group lasso projections:

$$\begin{aligned} (\mathbf{x}, \lambda, \eta) &= \text{SGLP}(\mathbf{v}, s_1, s_2), \\ (\mathbf{x}', \lambda', \eta') &= \text{SGLP}(\mathbf{v}, s'_1, s'_2). \end{aligned}$$

The following key result holds.

Theorem 3. *If $\lambda \leq \lambda'$ and $s_2 = s'_2$, then $s_1 \geq s'_1$.*

Theorem 3 gives the oracle procedure with its proof presented in the supplement. For a given estimator $\hat{\lambda}$, we compute its corresponding $\hat{\eta}$ from (12) and then \hat{s}_1 from (11), satisfying $(\hat{\mathbf{x}}, \hat{\lambda}, \hat{\eta}) = \text{SGLP}(\mathbf{v}, \hat{s}_1, s_2)$. Then \hat{s}_1 is compared with s_1 . Clearly, by Theorem 3, if $\hat{s}_1 \leq s_1$, the estimator $\hat{\lambda}$ is no less than λ^* . Otherwise, $\hat{s}_1 > s_1$ means $\hat{\lambda} < \lambda^*$. In addition, from (11) we know that \hat{s}_1 is a continuous function of $\hat{\lambda}$. Together with the monotonicity given in Theorem 3, a bisection approach can be employed to calculate λ^* . Algorithm 2 gives a detailed description.

4.3 Solving Restricted version of (7)

Finally, we modify the above procedures to compute the optimal solution of the restricted problem (6). To apply the accelerated gradient method, we consider the following projection step:

$$\begin{aligned} &\underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \\ &\text{subject to} && \|\mathbf{x}^{T_1}\|_1 \leq s_1 && (C_1) \\ &&& \|\mathbf{x}^{T_3}\|_G \leq s_2 && (C_2). \end{aligned} \tag{13}$$

Our first observation is: $T_3(\mathbf{x}) \subset T_1(\mathbf{x})$, since if an element of \mathbf{x} lies in a group whose L_2 -norm is less than τ , then the absolute value of this element must also be less than τ . Secondly, from the decomposable nature of the objective function, we conclude that:

$$x_j^* = \begin{cases} v_j & \text{if } j \in (T_1)^c \\ v_j^{\lambda^*} & \text{if } j \in T_1 \setminus T_3, \end{cases}$$

since there are no constraints on x_j if it is outside T_1 and involves only the L_1 -norm constraint if $j \in T_1 \setminus T_3$. Following routine calculations as in [9], we obtain the following results similar to (11) and (12):

$$s_1 = \sum_{i \in T_2} \max\{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2 - \eta^*, 0\} \frac{\|\mathbf{v}_{G_i}^{\lambda^*}\|_1}{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2} + \sum_{j \in T_1 \setminus T_3} v_j^{\lambda^*} \tag{14}$$

$$s_2 = \sum_{i \in T_2} \max\{\|\mathbf{v}_{G_i}^{\lambda^*}\|_2 - \eta^*, 0\}. \tag{15}$$

Based on (14) and (15), we design a similar bisection approach to compute λ^* and thus $(\mathbf{x}^*)^{T_3}$, as in Algorithm 2. Details are deferred to the supplement.

Algorithm 2 Sparse Group Lasso Projection Algorithm

Input: v, s_1, s_2 **Output:** an optimal solution x to the Sparse Group Projection Problem**Function** SGLP(v, s_1, s_2)

```
1: if  $\|x\|_1 \leq s_1$  and  $\|x\|_G \leq s_2$  then
2:   return  $v$ 
3: end if
4:  $x_{C_1} = \mathcal{P}_1^{s_1}(v)$ 
5:  $x_{C_2} = \mathcal{P}_G^{s_2}(v)$ 
6:  $x_{C_{12}} = \text{bisec}(v, s_1, s_2)$ 
7: if  $\|x_{C_1}\|_G \leq s_2$  then
8:   return  $x_{C_1}$ 
9: else if  $\|x_{C_2}\|_1 \leq s_1$  then
10:  return  $x_{C_2}$ 
11: else
12:  return  $x_{C_{12}}$ 
13: end if
```

Function bisec(v, s_1, s_2)

```
1: Initialize  $up, low$  and  $tol$ 
2: while  $up - low > tol$  do
3:    $\hat{\lambda} = (low + up)/2$ 
4:   if (12) has a solution  $\hat{\eta}$  given  $v^{\hat{\lambda}}$  then
5:     calculate  $\hat{s}_1$  using  $\hat{\eta}$  and  $\hat{\lambda}$ .
6:     if  $\hat{s}_1 \leq s_1$  then
7:        $up = \hat{\lambda}$ 
8:     else
9:        $low = \hat{\lambda}$ 
10:    end if
11:  else
12:     $up = \hat{\lambda}$ 
13:  end if
14: end while
15:  $\lambda^* = up$ 
16: Solve (12) to get  $\eta^*$ 
17: Calculate  $x^*$  from  $\lambda^*$  and  $\eta^*$  via (10)
18: return  $x^*$ 
```

5 Significance

This section is devoted to a brief discussion of advantages of our work statistically and computationally. Moreover, it explains why the proposed method is useful to perform efficient and interpretable feature selection with a given natural group structure.

Interpretability. The parameters in (2) are highly interpretable in that s_1 and s_2 are upper bounds of the number of nonzero elements as well as that of groups. This is advantageous, especially in the presence of certain prior knowledge regarding the number of features and/or that of groups. However, such an interpretation vanishes with convex methods such as lasso or sparse group lasso, in which incorporating such prior knowledge often requires repeated trials of different parameters.

Parameter tuning. Typically, tuning parameters for good generalization usually requires considerable amount work due to a large number of choices of parameters. However, tuning in (1) may search through integer values in a bounded range, and can be further simplified when certain prior knowledge is available.

Table 1: Running time (in seconds) of Dykstra’s, ADMM and our projection algorithm. All three algorithms are averaged over 100 replications.

METHODS	10^2	10^3	10^4	10^5	10^6
DYKSTRA	0.1944	0.5894	4.8702	51.756	642.60
ADMM	0.0519	0.1098	1.2000	26.240	633.00
OURS	$< 10^{-7}$	0.0002	0.0051	0.0440	0.5827

This permits more efficient tuning than its regularization counterpart. Based on our limited experience, we note that τ does not need to be tuned precisely as we may fix at some small values.

Performance and Computation. Although our model (2) is proposed as a computational surrogate of the ideal L_0 -method, its performance can also be theoretically guaranteed, i.e., consistent feature selection can be achieved. Moreover, the computation of our model is much more efficient and applicable to large-scale applications.

6 Empirical Evaluation

This section performs numerical experiments to evaluate the proposed methods in terms of the efficiency and accuracy of sparse group feature selection. Evaluations are conducted on a PC with i7-2600 CPU, 8.0 GB memory and 64-bit Windows operating system.

6.1 Evaluation of Projection Algorithms

Since the DC programming and the accelerated gradient methods are both standard, the efficiency of the proposed nonconvex formulation (2) depends on the projection step in (8). Therefore, we focus on evaluating the projection algorithms and comparing with two popular projection algorithms: Alternating Direction Multiplier Method (ADMM) [5] and Dykstra’s projection algorithm [7]. We provide a detailed derivation of adapting these two algorithms to our formulation in the supplement.

To evaluate the efficiency, we first generate the vector \mathbf{v} whose entries are uniformly distributed in $[-50, 50]$ and the dimension of \mathbf{v} , denoted as p , is chosen from the set $\{10^2, 10^3, 10^4, 10^5, 10^6\}$. Next we partition the vector into 10 groups of equal size. Finally, s_2 is set to $5 \log(p)$ and s_1 , the radius of the L_1 -ball, is computed by $\frac{\sqrt{10}}{2}s_2$ (motivated by the fact that $s_1 \leq \sqrt{10}s_2$).

For a fair comparison, we run our projection algorithm until converge and record the minimal objective value as f^* . Then we run ADMM and Dykstra’s algorithm until their objective values become close to ours. More specifically, we terminate their iterations as soon as $f_{\text{ADMM}} - f^* \leq 10^{-3}$ and $f_{\text{Dykstra}} - f^* \leq 10^{-3}$, where f_{ADMM} and f_{Dykstra} stand for the objective value of ADMM and Dykstra’s algorithm respectively. Table 1 summarizes the average running time of all three algorithms over 100 replications.

Next we demonstrate the accuracy of our projection algorithm. Toward this end, the general convex optimization toolbox CVX [12] is chosen as the baseline. Following the same strategy of generating data, we report the distance (computed from the Euclidean norm $\|\cdot\|_2$) between optimal solution of the three projection algorithms and that of the CVX. Note that the projection is strictly convex with a unique global optimal solution.

For ADMM and Dykstra’s algorithm, the termination criterion is that the relative difference of the objective values between consecutive iterations is less than a threshold value. Specifically, we terminate the iteration if $|f(\mathbf{x}_{k-1}) - f(\mathbf{x}_k)| \leq 10^{-7}f(\mathbf{x}_{k-1})$. For our projection algorithm, we set the tol in Algorithm 2 to be 10^{-7} . The results are summarized in Table 2. Powered by second-order optimization algorithms, CVX can

provide fast and accurate solution for problems of moderate size but would suffer from great computational burden for large-scale ones. Therefore we only report the results up to 5,000 dimensions.

Table 2: Distance between the optimal solution of projection algorithms and that of the CVX. All the results are averaged over 100 replications.

Methods	50	100	500	1000	5000
Dykstra	9.00	9.81	11.40	11.90	12.42
ADMM	0.64	0.08	3.6e-3	6.3e-3	1.3e-2
ours	1.4e-3	1.1e-3	1.2e-3	1.7e-3	7.3e-3

From Tables 1 and 2, we note that both ADMM and our algorithm yield more accurate solution than that of Dykstra’s. For projections of moderate size, all three algorithms perform well. However, for large-scale ones, our advantage on efficiency is evident.

6.2 Performance on Synthetic Data

6.2.1 Experimental Setup

We generate a 60×100 matrix \mathbf{A} , whose entries follow i.i.d standard normal distribution. The 100 features (columns) are partitioned into 10 groups of equal size. The ground truth vector \mathbf{x}_0 possesses nonzero elements only in 4 of the 10 groups. To further enhance sparsity, in each nonzero group of \mathbf{x}_0 , only t ($t \leq 10$) elements are nonzero, where t is uniformly distributed from $[1, 5]$. Finally \mathbf{y} is generated according to $\mathbf{A}\mathbf{x}_0 + \mathbf{z}$ with \mathbf{z} following distribution $\mathcal{N}(0, 0.5^2)$, where \mathbf{A} and \mathbf{y} are divided into training and testing set of equal size.

We fit our method to the training set and compare with lasso, group lasso and sparse group lasso. The tuning parameters of the convex methods are selected from $\{0.01, 0.1, 1, 10\}$, whereas for our method, the number of nonzero groups is selected from the set $\{2, 4, 6, 8\}$ and the number of features is chosen from $\{2s_2, 4s_2, 6s_2, 8s_2\}$. Leave-one-out cross-validation is conducted over the training set for choosing the best tuning parameter for all the methods.

6.2.2 Results and Discussions

We use following metrics for evaluation:

- Estimation error: $\|\hat{\mathbf{x}} - \mathbf{x}_0\|_2^2$
- Prediction error: $\|\mathbf{A}\hat{\mathbf{x}} - \tilde{\mathbf{y}}\|_2^2$
- Group precision: $|T_2(\hat{\mathbf{x}}) \cap T_2(\mathbf{x}_0)|/|T_2(\hat{\mathbf{x}})|$
- Group recall: $|T_2(\hat{\mathbf{x}}) \cap T_2(\mathbf{x}_0)|/|T_2(\mathbf{x}_0)|$

where $\hat{\mathbf{x}}$ is the estimator obtained from (2) and $\tilde{\mathbf{y}}$ is an independent vector following the same distribution as \mathbf{y} . The group precision and recall demonstrate the capability of recovering the group structure from data. We report the results in Table 3 and observe that our model generally exhibits better performance. Note that although our model does not provide the best result on the metric of group recall, the group precision of our model is significantly better than the others, illustrating the fact that the three convex methods recover more redundant groups.

Table 3: Comparison of Performance on synthetic data, where glasso stands for the group lasso and sglasso denotes sparse group lasso. All the results are averaged for 100 replications.

METHODS	ESTI.	PRED.	PREC.	REC.
LASSO	4.7933	151.05	0.5212	0.8700
GLASSO	8.1230	244.53	0.5843	0.7575
SGLASSO	4.7649	151.29	0.5215	0.8675
OURS	4.6617	142.18	0.7848	0.6450

6.3 Performance on Real-world Application

Our method is further evaluated on the application of examining Electroencephalography (EEG) correlates of genetic predisposition to alcoholism [10]. EEG records the brain’s spontaneous electrical activity by measuring the voltage fluctuations over multiple electrodes placed on the scalp. This technology has been widely used in clinical diagnosis, such as coma, brain death and genetic predisposition to alcoholism. In fact, encoded in the EEG data is a certain group structure, since each electrode records the electrical activity of a certain region of the scalp. Identifying and utilizing such spatial information has the potential of increasing stability of a prediction.

The training set contains 200 samples of 16384 dimensions, sampled from 64 electrodes placed on subject’s scalps at 256 Hz (3.9-msec epoch) for 1 second. Therefore, the data can naturally be divided into 64 groups of size 256. We apply the lasso, group lasso, sparse group lasso and our method on the training set and adapt the 5-fold cross-validation for selecting tuning parameters. More specifically, for lasso and group lasso, the candidate tuning parameters are specified by 10 parameters³ sampled using the logarithmic scale from the parameter spaces, while for the sparse group lasso, the parameters form a 10×10 grid⁴, sampled from the parameter space in logarithmic scale. For our method, the number of groups is selected from the set: $s_2 = \{30, 40, 50\}$ and s_1 , the number of features is chosen from the set $\{50s_2, 100s_2, 150s_2\}$. The accuracy of classification together with the number of selected features and groups over a test set, which also contains 200 samples, are reported in Table 4. Clearly our method achieves the best performance of classification with the least number of groups. Note that, although lasso’s performance is almost as good as ours with even less features, however, it fails to identify the underlying group structure in the data, as revealed by the fact all 64 groups are selected.

Table 4: Comparison of performance on EEG data, where glasso stands for group lasso and sglasso denotes sparse group lasso.

Methods	Accuracy	# Feature	# Group
lasso	67.0	2068	64
glasso	62.5	8704	34
sglasso	65.5	4834	61
ours	68.0	3890	25

7 Conclusion and Future Work

This paper expands a nonconvex paradigm into sparse group feature selection. In particular, theoretical properties on the accuracy of selection and parameter estimation are analyzed. In addition, an efficient optimization scheme is developed based on the DC programming, accelerated gradient method and efficient

³ $\lambda_{\text{lasso}} = \text{logspace}(10^{-3}, 1)$, $\lambda_{\text{glasso}} = \text{logspace}(10^{-2}, 1)$

⁴The product space of $\lambda_{\text{lasso}} \times \lambda_{\text{glasso}}$

projection. The efficiency and efficacy of the proposed method are validated on both synthetic data and real-world applications.

The proposed method will be further investigated on real-world applications involving the group structure. Moreover, extending the proposed model to multi-modal multi-task learning [25] is another promising direction.

References

- [1] Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. *Convex Optimization with Sparsity-Inducing Norms*. 2010.
- [2] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [3] P.J. Bickel, Y. Ritov, and A.B. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- [4] J.F. Bonnans and A. Shapiro. Optimization problems with perturbations: A guided tour. *SIAM review*, pages 228–264, 1998.
- [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers, 2011.
- [6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [7] P.L. Combettes and J.C. Pesquet. Proximal splitting methods in signal processing. *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, 2010.
- [8] D.L. Donoho. De-noising by soft-thresholding. *Information Theory, IEEE Transactions on*, 41(3):613–627, 2002.
- [9] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning*, pages 272–279. ACM, 2008.
- [10] A. Frank and A. Asuncion. UCI machine learning repository, 2010.
- [11] J. Friedman, T. Hastie, and R. Tibshirani. A note on the group lasso and a sparse group lasso. *Arxiv preprint arXiv:1001.0736*, 2010.
- [12] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, April 2011.
- [13] J. Huang and T. Zhang. The benefit of group sparsity. *The Annals of Statistics*, 38(4):1978–2004, 2010.
- [14] A.N. Kolmogorov and V.M. Tihomirov. *e-Entropy and e-capacity of sets in functional spaces*. American Mathematical Society, 1961.
- [15] R. Mazumder, J.H. Friedman, and T. Hastie. Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association*, 106(495):1125–1138, 2011.
- [16] L. Meier, S. Van De Geer, and P. Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- [17] B. K. Natarajan. Sparse approximation solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995.
- [18] Y. Nesterov. Gradient methods for minimizing composite objective function. *CORE Discussion Papers*, 2007.
- [19] X. Shen, W. Pan, and Y. Zhu. Likelihood-based selection and sharp parameter estimation. *Journal of American Statistical Association*, to appear, 2012.
- [20] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [21] S.A. Van De Geer and P. Bühlmann. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- [22] W.H. Wong and X. Shen. Probability inequalities for likelihood ratios and convergence rates of sieve mles. *The Annals of Statistics*, 23(2):339–362, 1995.

- [23] H. Yang, Z. Xu, I. King, and M. Lyu. Online learning for group lasso. In *Proceedings of the 27th International Conference on Machine Learning*, pages 1191–1198. ACM, 2010.
- [24] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [25] D. Zhang and D. Shen. Multi-modal multi-task learning for joint prediction of multiple regression and classification variables in alzheimer’s disease. *NeuroImage*, 2011.
- [26] T. Zhang. Analysis of multi-stage convex relaxation for sparse regularization. *The Journal of Machine Learning Research*, 11:1081–1107, 2010.
- [27] T. Zhang. Multi-stage convex relaxation for feature selection. *arXiv:1106.0565v2*, 2011.
- [28] P. Zhao and B. Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.

Supplementary Material for paper: Efficient Sparse Group Feature Selection via Nonconvex Optimization

1 Proof of Theorem 1

The proof uses a large deviation probability inequality of [22] to treat one-sided log-likelihood ratios with constraints.

Let $\mathcal{S} = \{\mathbf{x}^\tau : \|\mathbf{x}^\tau\|_0 \leq s_1^0, \|\mathbf{x}^\tau\|_{0,G} \leq s_2^0\}$, $\|\mathbf{x}\|_0 = \sum_{j=1}^p I(|x_j| \neq 0)$ is the L_0 -norm of \mathbf{x} , and $\|\mathbf{x}\|_{0,G} = \sum_{j=1}^{|G|} I(\|\mathbf{x}_j\|_2 \neq 0)$ is the L_0 -norm over the groups. Now we partition \mathcal{S} . Note that for $C \subset (G_1, \dots, G_{|G|})$, it can be partitioned into $C = (C \setminus C^0) \cup (C \cap C^0)$. Then

$$\mathcal{S} = \bigcup_{i=0}^{s_2^0} \bigcup_{C \in \mathcal{B}_i} \mathcal{S}_{A_C, C},$$

where $\mathcal{S}_{A_C, C} = \{\mathbf{x}^\tau \in \mathcal{S} : C(\mathbf{x}) = C = (G_{i_1}, \dots, G_{i_k}), \sum_j |A_{G_j}| \leq s_1^0\}$, and $\mathcal{B}_i = \{C \neq C_0 : |C^0 \setminus C| = i, |C| \leq s_2^0\}$, with $|\mathcal{B}_i| = \binom{s_2^0}{s_2^0 - i} \sum_{j=0}^i \binom{|G| - s_2^0}{j}$; $i = 0, \dots, s_2^0$.

To bound the error probability, let $L(\mathbf{x}) = -\frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2$ be the likelihood. Note that

$$\{\hat{\mathbf{x}} \neq \hat{\mathbf{x}}^o\} \subseteq \{L(\hat{\mathbf{x}}) - L(\hat{\mathbf{x}}^o) \geq 0\} \subseteq \{L(\hat{\mathbf{x}}) - L(\mathbf{x}^0) \geq 0\}.$$

This together with $\{\hat{\mathbf{x}} \neq \hat{\mathbf{x}}^o\} \subseteq \{\hat{\mathbf{x}} \in \mathcal{S}\}$ implies that

$$\{\hat{\mathbf{x}} \neq \hat{\mathbf{x}}^o\} \subseteq \{L(\hat{\mathbf{x}}) - L(\mathbf{x}^0) \geq 0\} \cap \{\hat{\mathbf{x}} \in \mathcal{S}\}.$$

Consequently,

$$\begin{aligned} I &\equiv P(\hat{\mathbf{x}} \neq \hat{\mathbf{x}}^o) \\ &\leq P\left(L(\hat{\mathbf{x}}) - L(\mathbf{x}^0) \geq 0; \hat{\mathbf{x}} \in \mathcal{S}\right) \\ &\leq \sum_{i=1}^{s_2^0} \sum_{C \in \mathcal{B}_i} \sum_{\mathcal{S}_{A_C, C}} P^*\left(\sup_{\mathbf{x} \in \mathcal{S}_{A_C, C}} (L(\mathbf{x}) - L(\mathbf{x}^0)) \geq 0\right) \\ &\leq \sum_{i=1}^{s_2^0} \sum_{j=1}^{s_1^0} \sum_{|C|=i, |A_G|=j} P^*\left(\sup_{\left\{-\log(1-h^2(\mathbf{x}, \mathbf{x}^0)) \geq \max(i, 1)C_{\min}(\mathbf{x}^0) - d_3\tau^{d_2}p, \mathbf{x} \in \mathcal{S}_{A_C, C}\right\}} (L(\mathbf{x}) - L(\mathbf{x}^0)) \geq 0\right), \end{aligned}$$

where P^* is the outer measure and the last two inequalities use the fact that $\mathcal{S}_{A_C, C} \subseteq \{\mathbf{x} \in \mathcal{S}_{A_C, C} : \max(|C^0 \setminus C|, 1)C_{\min}(\mathbf{x}^0) \leq -\log(1-h^2(\mathbf{x}, \mathbf{x}^0))\} \subseteq \{-\log(1-h^2(\mathbf{x}, \mathbf{x}^0)) \geq d_1 \max(i, 1)C_{\min}(\mathbf{x}^0) - d_3\tau^{d_2}p\}$, under Assumption 3.

For I , we apply Theorem 1 of [22] to bound each term. Towards this end, we verify their entropy condition (3.1) for the local entropy over $\mathcal{S}_{A_C, C}$ for $|C| = 1, \dots, s_2^0$ and $|A| = 1, \dots, s_1^0$. Under Assumption 2 $\varepsilon = \varepsilon_{n,p} = (2c_0)^{1/2}c_4^{-1} \log(2^{1/2}/c_3) \log p (\frac{s_1^0}{n})^{1/2}$ satisfies there with respect to $\varepsilon > 0$, that is,

$$\sup_{\{0 \leq |A| \leq p_0\}} \int_{2^{-8\varepsilon^2}}^{2^{1/2}\varepsilon} H^{1/2}(t/c_3, \mathcal{F}_{ji}) dt \leq p_0^{1/2} 2^{1/2} \varepsilon \log(2/2^{1/2}c_3) \leq c_4 n^{1/2} \varepsilon^2. \quad (16)$$

for some constant $c_3 > 0$ and $c_4 > 0$, say $c_3 = 10$ and $c_4 = \frac{(2/3)^{5/2}}{512}$. By Assumption 2, $C_{\min}(\mathbf{x}^0) \geq \varepsilon_{n,p_0,p}^2$ implies (16), provided that $s_1^0 \geq (2c_0)^{1/2} c_4^{-1} \log(2^{1/2}/c_3)$.

Note that $|\mathcal{B}_i| = \binom{s_2^0}{s_2^0-i} \sum_{j=0}^i \binom{|G|-s_2^0}{j} \leq (|G|(|G|-s_2^0))^i \leq (|G|^2/4)^i$ by the binomial coefficients formula. Moreover, $\sum_{j=1}^{s_1^0} 2^j i^j \leq i^{s_1^0}$, and $\sum_{j_1+\dots+j_i=j} \binom{j}{j_1, \dots, j_i} 2^j = (2i)^j$ using the Multinomial Theorem. By Theorem 1 of [22], there exists a constant $c_2 > 0$, say $c_2 = \frac{4}{27} \frac{1}{1926}$,

$$\begin{aligned} I &\leq \sum_{i=1}^{s_2^0} |\mathcal{B}_i| \sum_{j=1}^{s_1^0} \sum_{(j_1, \dots, j_i)} \binom{j}{j_1, \dots, j_i} 2^{j_1} \dots 2^{j_i} \exp(-c_2 n i C_{\min}(\mathbf{x}^0)) \\ &\leq \sum_{i=1}^{s_2^0} \exp(-c_2 n i C_{\min}(\mathbf{x}^0) + 2i(\log |G| + \log s_1^0)) \\ &\leq \exp(-c_2 n C_{\min}(\mathbf{x}^0) + 2(\log |G| + \log s_1^0)). \end{aligned}$$

Let $G = \{\hat{\mathbf{x}} \neq \hat{\mathbf{x}}^0\}$. For the risk property, $Eh^2(\hat{\mathbf{x}}, \mathbf{x}^0) = Eh^2(\hat{\mathbf{x}}^0, \mathbf{x}^0) + Eh^2(\hat{\mathbf{x}}, \mathbf{x}^0)I(G)$ is upper bounded by

$$Eh^2(\hat{\mathbf{x}}, \mathbf{x}^0) + \exp(-c_2 n C_{\min}(\mathbf{x}^0) + 2(\log |G| + \log s_1^0)) = (1 + o(1))Eh^2(\hat{\mathbf{x}}^0, \mathbf{x}^0),$$

using the fact that $h(\hat{\mathbf{x}}, \mathbf{x}^0) \leq 1$. This completes the proof.

2 Proof of Theorem 3

We utilize an intermediate lemma from [4]:

Lemma 2. *Let X be a metric space and U be a normed space. Suppose that for all $x \in X$, the function $\psi(x, \cdot)$ is differentiable and that $\psi(x, Y)$ and $D_Y \psi(x, Y)$ (the partial derivative of $\psi(x, Y)$ with respect to Y) are continuous on $X \times U$. Let Φ be a compact subset of X . Define the optimal value function as $\phi(Y) = \inf_{x \in \Phi} \psi(x, Y)$. The optimal value function $\phi(Y)$ is directionally differentiable. In addition, if for any $Y \in U$, $\psi(\cdot, Y)$ has a unique minimizer $x(Y)$ over Φ , then $\phi(Y)$ is differentiable at Y and the gradient of $\phi(Y)$ is given by $\phi'(Y) = D_Y \psi(x(Y), Y)$.*

Proof of Theorem 3. For the proof, an intermediate lemma will be used, with its details given in the Appendix. Since both constraints are active, if $(x, \lambda, \eta) = \text{SGLP}(v, s_1, s_2)$, then x and λ are also the optimal solutions to the following problem:

$$\underset{\lambda}{\text{maximize}} \underset{x \in X}{\text{minimize}} \quad \psi(x, \lambda) = \frac{1}{2} \|x - v\|_2^2 + \lambda(\|x\|_1 - s_1),$$

where $X = \{x : \|x\|_G \leq s_2\}$. By Lemma 2, $\phi(\lambda) = \inf_{x \in X} \psi(x, \lambda)$ is differentiable with the derivative given by $\|x\|_1$. In addition, as a pointwise infimum of a concave function, so does $\phi(\lambda)$ [6] and its derivative, $\|x\|_1$, is non-increasing. Therefore $s_1 = \|x\|_1$ is non-decreasing as λ becomes smaller. This completes the proof. \square

3 Algorithm for Solving (13)

We give a detailed description of algorithm for solving the restricted projection (13) in Algorithm 3.

Algorithm 3 Restricted Sparse Group Lasso Projection Algorithm

Input: \mathbf{v} , s_1 , s_2 , T_1 , T_3 **Output:** an optimal solution \mathbf{x} to the Restricted Sparse Group Projection Problem (13)**Function** RSGLP(\mathbf{v} , s_1 , s_2 , T_1 , T_3)

```
1: if  $\|\mathbf{x}^{T_1}\|_1 \leq s_1$  and  $\|\mathbf{x}^{T_3}\|_G \leq s_2$  then
2:   return  $\mathbf{v}$ 
3: end if
4:  $\mathbf{x}_{C_1}^{(T_1)^c} = \mathbf{v}^{(T_1)^c}$ ,  $\mathbf{x}_{C_1}^{T_1} = \mathcal{P}_1^{s_1}(\mathbf{v}^{T_1})$ 
5:  $\mathbf{x}_{C_2}^{(T_3)^c} = \mathbf{v}^{(T_3)^c}$ ,  $\mathbf{x}_{C_2}^{T_3} = \mathcal{P}_G^{s_2}(\mathbf{v}^{T_3})$ 
6:  $\mathbf{x}_{C_{12}}^{(T_1)^c} = \mathbf{v}^{(T_1)^c}$ ,  $\mathbf{x}_{C_{12}}^{T_1} = \text{bisec}(\mathbf{v}, s_1, s_2, T_1, T_3)$ 
7: if  $\|\mathbf{x}_{C_1}^{T_3}\|_G \leq s_2$  then
8:   return  $\mathbf{x}_{C_1}$ 
9: else if  $\|\mathbf{x}_{C_2}^{T_1}\|_1 \leq s_1$  then
10:  return  $\mathbf{x}_{C_2}$ 
11: else
12:  return  $\mathbf{x}_{C_{12}}$ 
13: end if
```

Function bisec(\mathbf{v} , s_1 , s_2 , T_1 , T_3)

```
1: Initialize  $up$ ,  $low$  and  $tol$ 
2: while  $up - low > tol$  do
3:    $\hat{\lambda} = (low + up)/2$ 
4:   if (15) has a solution  $\hat{\eta}$  given  $v^{\hat{\lambda}}$  then
5:     calculate  $\hat{s}_1$  using  $\hat{\eta}$  and  $\hat{\lambda}$ .
6:     if  $\hat{s}_1 \leq s_1$  then
7:        $up = \hat{\lambda}$ 
8:     else
9:        $low = \hat{\lambda}$ 
10:    end if
11:  else
12:     $up = \hat{\lambda}$ 
13:  end if
14: end while
15:  $\lambda^* = up$ 
16: Solve (15) to get  $\eta^*$ 
17: Calculate  $(\mathbf{x}^*)^{T_1}$  from  $\lambda^*$  and  $\eta^*$ .
18: return  $(\mathbf{x}^*)^{T_1}$ 
```

4 Accelerated Gradient Method

The AGM procedure is listed in Algorithms 4, in which $f(\mathbf{x})$ is the objective function $\frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$ with $\nabla f(\mathbf{x})$ denotes its gradient at \mathbf{x} . In addition, $f_{L,\mathbf{u}}(\mathbf{x})$ is the linearization of $f(\mathbf{x})$ at \mathbf{u} defined as follows:

$$f_{L,\mathbf{u}}(\mathbf{x}) = f(\mathbf{u}) + \nabla f(\mathbf{u})^T(\mathbf{x} - \mathbf{u}) + \frac{L}{2}\|\mathbf{x} - \mathbf{u}\|_2^2.$$

Algorithm 4 Accelerated Gradient Method [18, 2] for (7)

Input: $\mathbf{A}, \mathbf{y}, s_1, s_2, L_0, \mathbf{x}_0$,

Output: solution \mathbf{x} to (7)

1: **Initialize:** $L_0, \mathbf{x}_1 = \mathbf{x}_0, \alpha_{-1} = 0, \alpha_0 = 1, t = 0$.

2: **repeat**

3: $t = t + 1, \beta_t = \frac{\alpha_{t-2}-1}{\alpha_{t-1}}, \mathbf{u}_t = \mathbf{x}_t + \beta_t(\mathbf{x}_t - \mathbf{x}_{t-1})$

4: **Line search:** Find the smallest $L = 2^j L_{t-1}$ such that

$$f(\mathbf{x}_{t+1}) \leq f_{L, \mathbf{u}_t}(\mathbf{x}_{t+1}),$$

 where $\mathbf{x}_{t+1} = \text{SGLP}(\mathbf{u}_t - \frac{1}{L} \nabla f(\mathbf{u}_t), s_1, s_2)$

5: $\alpha_{t+1} = \frac{1 + \sqrt{1 + 4\alpha_t^2}}{2}, L_t = L$.

6: **until** Converge

7: **return** \mathbf{x}_t

5 ADMM Projection algorithm

ADMM is widely chosen for its capability of decomposing coupled variables/constraints, which is exactly the case in our projection problem. Before applying ADMM, we transform (8) into an equivalent form as follows:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \\ & \text{subject to} && \|\mathbf{u}\|_1 \leq s_1 \\ & && \|\mathbf{w}\|_G \leq s_2 \\ & && \mathbf{u} = \mathbf{x}, \mathbf{w} = \mathbf{x}. \end{aligned}$$

The augmented Lagrangian is:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\eta}) = & \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + \boldsymbol{\lambda}^T (\mathbf{u} - \mathbf{x}) + \boldsymbol{\eta}^T (\mathbf{w} - \mathbf{x}) \\ & + \frac{\rho}{2} (\|\mathbf{u} - \mathbf{x}\|_2^2 + \|\mathbf{w} - \mathbf{x}\|_2^2). \end{aligned}$$

Utilize the scaled form [5], i.e., let $\boldsymbol{\lambda} = \frac{\boldsymbol{\lambda}}{\rho}, \boldsymbol{\eta} = \frac{\boldsymbol{\eta}}{\rho}$, we can obtain an equivalent augmented Lagrangian:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\eta}) = & \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + \frac{\rho}{2} (\|\mathbf{x} - \mathbf{u} - \boldsymbol{\lambda}\|_2^2 + \|\mathbf{x} - \mathbf{w} - \boldsymbol{\eta}\|_2^2) \\ & - \frac{\rho}{2} (\|\boldsymbol{\lambda}\|_2^2 + \|\boldsymbol{\eta}\|_2^2). \end{aligned}$$

Now we calculate the optimal $\mathbf{x}, \boldsymbol{\lambda}$ and $\boldsymbol{\eta}$ through alternating minimization. For fixed \mathbf{u} and \mathbf{w} , the optimal \mathbf{x} possesses a closed-form solution:

$$\mathbf{x} = \frac{1}{1 + 2\rho} (\mathbf{v} + \rho(\mathbf{u} + \boldsymbol{\lambda} + \mathbf{w} + \boldsymbol{\eta})).$$

For fixed \mathbf{x} and \mathbf{u} , finding the optimal \mathbf{w} is a group lasso projection:

$$\begin{aligned} & \underset{\mathbf{w}}{\text{minimize}} && \frac{1}{2} \|\mathbf{w} - (\mathbf{x} - \boldsymbol{\eta})\|_2^2 \\ & \text{subject to} && \|\mathbf{w}\|_G \leq s_2 \end{aligned} \tag{17}$$

For fixed \mathbf{x} and \mathbf{w} , finding the optimal \mathbf{u} amounts to solve an L_1 -ball projection:

$$\begin{aligned} & \underset{\mathbf{u}}{\text{minimize}} && \frac{1}{2} \|\mathbf{u} - (\mathbf{x} - \boldsymbol{\lambda})\|_2^2 \\ & \text{subject to} && \|\mathbf{u}\|_1 \leq s_1. \end{aligned} \tag{18}$$

The update of multipliers is standard as follows:

$$\begin{aligned} \boldsymbol{\lambda} &= \boldsymbol{\lambda} + \mathbf{u} - \mathbf{x} \\ \boldsymbol{\eta} &= \boldsymbol{\eta} + \mathbf{w} - \mathbf{x} \end{aligned} \tag{19}$$

Algorithm 5 summarizes the above procedure. Note that, the value of the penalty term ρ is fixed in Algorithm 5. However, in our implementation, we increase ρ whenever necessary to obtain faster convergence.

Algorithm 5 ADMM [5] for (8)

Input: \mathbf{v}, s_1, s_2

Output: an optimal solution \mathbf{x} to (8)

- 1: **Initialize:** $\mathbf{x}_0, \mathbf{u}_0, \mathbf{w}_0, \boldsymbol{\lambda}_0, \boldsymbol{\eta}_0, t = 0, \rho > 0$
 - 2: **repeat**
 - 3: $t = t + 1$
 - 4: $\mathbf{x}_t = \frac{1}{1+2\rho} (\mathbf{v} + \rho(\mathbf{u}_{t-1} + \boldsymbol{\lambda}_{t-1} + \mathbf{w}_{t-1} + \boldsymbol{\eta}_{t-1}))$
 - 5: $\mathbf{w}_t = \mathcal{P}_G^{s_2}(\mathbf{x}_t - \boldsymbol{\eta}_{t-1})$
 - 6: $\mathbf{u}_t = \mathcal{P}_1^{s_1}(\mathbf{x}_t - \boldsymbol{\lambda}_{t-1})$
 - 7: $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}_{t-1} + \mathbf{u}_t - \mathbf{x}_t, \boldsymbol{\eta}_t = \boldsymbol{\eta}_{t-1} + \mathbf{w}_t - \mathbf{x}_t.$
 - 8: **until** Converge
 - 9: **return** \mathbf{x}_t
-

6 Dykstra's Algorithm

Dykstra's algorithm is a general scheme to compute the projection onto intersections of convex sets. It is carried out by taking Euclidean projections onto each convex set alternatively in a smart way and is guaranteed to converge for least squares objective function [7]. The details of applying Dykstra's Algorithm to our projection problem are listed in Algorithm 6.

Algorithm 6 Dykstra's Algorithm [7] for (8)

Input: \mathbf{v}, s_1, s_2

Output: an optimal solution \mathbf{x} to (8)

- 1: **Initialize:** $\mathbf{x}_0 = \mathbf{v}, \mathbf{p}_0 = \mathbf{0}, \mathbf{q}_0 = \mathbf{0}, t = 0$
 - 2: **repeat**
 - 3: $t = t + 1$
 - 4: $\mathbf{y}_{t-1} = \mathcal{P}_G^{s_2}(\mathbf{x}_{t-1} + \mathbf{p}_{t-1})$
 - 5: $\mathbf{p}_t = \mathbf{x}_{t-1} + \mathbf{p}_{t-1} - \mathbf{y}_{t-1}$
 - 6: $\mathbf{x}_t = \mathcal{P}_1^{s_1}(\mathbf{y}_{t-1} + \mathbf{q}_{t-1})$
 - 7: $\mathbf{q}_t = \mathbf{y}_{t-1} + \mathbf{q}_{t-1} - \mathbf{x}_t$
 - 8: **until** Converge
 - 9: **return** \mathbf{x}_t
-